

Efficient Model Monitoring for Quality Control in Cardiac Image Segmentation^{*}

Francesco Galati¹ and Maria A. Zuluaga¹[0000-0002-1147-766X]

Data Science Department
EURECOM, Sophia Antipolis, France
{galati,zuluaga}@eurecom.fr

Abstract. Deep learning methods have reached state-of-the-art performance in cardiac image segmentation. Currently, the main bottleneck towards their effective translation into clinics requires assuring continuous high model performance and segmentation results. In this work, we present a novel learning framework to monitor the performance of heart segmentation models in the absence of ground truth. Formulated as an anomaly detection problem, the monitoring framework allows deriving surrogate quality measures for a segmentation and allows flagging suspicious results. We propose two different types of quality measures, a global score and a pixel-wise map. We demonstrate their use by reproducing the final rankings of a cardiac segmentation challenge in the absence of ground truth. Results show that our framework is accurate, fast, and scalable, confirming it is a viable option for quality control monitoring in clinical practice and large population studies.

Keywords: Quality Control · Cardiac Segmentation · Machine Learning

1 Introduction

With the advent of learning-based techniques over the last decade, cardiac image segmentation has reached state-of-the-art performance [5]. This achievement has opened the possibility to develop image segmentation frameworks that can assist and automate (partially or fully) the image analysis pipelines of large-scale population studies or routine clinical procedures.

The current bottleneck towards the large-scale use of learning-based pipelines in the clinic comes from the monitoring and maintenance of the deployed machine learning systems [15]. As shown in [5], despite the very high performances achieved, these methods may generate anatomically impossible results. In clinical practice and population studies, it is of utmost importance to constantly monitor a model’s performance to determine when it degrades or fails, leading to poor quality results, as they may represent important risks. A system’s continuous performance assessment and the detection of its degradation are challenging after deployment, due to the lack of a reference or ground truth. Therefore,

^{*} This work was supported through funding from the Monaco Government.

translation of models into clinical practice requires the development of monitoring mechanisms to measure a model’s segmentation quality, in the absence of ground truth, which guarantee their safe use in clinical routine and studies.

In a first attempt to assess performances of cardiac segmentation models in the absence of ground truth, Robinson et al. [12] proposed a supervised DL-based approach to predict the Dice Score Coefficient (DSC). More recently, Puyol-Antón et al. [11] used a Bayesian neural network to measure a model’s performance by classifying its resulting segmentation as correct or incorrect, whereas Ruijsink et al. [14] also use qualitative labels to train a support vector machine that predicts both the quality of the segmentation and of derived functional parameters. The main drawback of these methods is that they require annotations reflecting a large set of quantitative (e.g. DSC) or qualitative (e.g. correct/incorrect) segmentation quality levels, which can be difficult to obtain. The Reverse Classification Accuracy (RCA) [13,18] addresses this problem by using atlas label propagation. This registration-based method relies on the spatial overlap between the predicted segmentation and a reference atlas. It works under the hypothesis that if the predicted segmentation is of good quality, then it will produce a good segmentation on at least one atlas image. However, the atlas registration step may fail. This is often the case for certain cardiac pathologies that introduce significant morphological deformations that the registration step is not able to recover [21]. In such cases, it is necessary to verify the results and manually fine-tune the registration step, limiting the method’s scalability.

We present a novel learning framework to monitor the performance of cardiac image segmentation models in the absence of ground truth. The proposed framework is formulated as an anomaly detection problem. The intuition behind this work lies in the possibility of estimating a model of the variability of cardiac segmentation masks from a reference training dataset provided with reliable ground truth. This model is represented by a convolutional autoencoder, which can be subsequently used to identify anomalies in segmented unseen images. Differently from previous learning-based approaches [11,12], we avoid the need of any type of annotations about the quality of a segmentation for training. Our approach also avoids the required spatial alignment between ground truth images and segmentations of RCA [13,18], thus circumventing image registration.

2 Method

Let us denote $X \in \mathcal{C}^{H \times W}$ a segmentation mask of width W and height H , with \mathcal{C} the set of possible label values. A Convolutional Autoencoder (CA) is trained to learn a function $f : \mathcal{C}^{H \times W} \rightarrow \mathcal{C}^{H \times W}$, with $X' = f(X) \approx X$, by minimizing a global dissimilarity measure between an input mask X and its reconstruction X' . In an anomaly detection setup, the CA is trained using normal samples, *i.e.* samples without defects. In our framework, the normal samples are the ground truth (GT) masks associated with the images used to train a segmentation model (Fig. 1a). The CA learns to reconstruct defect-free samples, *i.e.* the GT, through a bottleneck, the latent space Z .




Fig. 1. **a.** A Convolutional Autoencoder (CA) is trained with ground truth (GT) masks from a cardiac imaging dataset. **b.** At inference, the CA reconstructs an input mask \hat{X} , previously segmented by a model. The reconstructed mask \hat{X}' acts as pseudo ground truth (pGT) to estimate a function $\rho(\hat{X}, \hat{X}')$, a surrogate measure of the segmentation quality and the model’s performance.

At inference (Fig. 1b), the CA is used to obtain $\hat{X}' = f(\hat{X})$, where \hat{X} is a segmentation mask, generated by a cardiac segmentation model/method on unseen data, and \hat{X}' its reconstruction. Since the CA is trained with ground truth data, the quality of the reconstruction will be generally higher for segmentation masks with similar characteristics than those in the ground truth. Poor segmentations, which the CA has not encountered at training, will instead lead to bad reconstructions ($\hat{X}' \not\approx \hat{X}$). Autoencoder-based anomaly detection methods exploit the reconstruction error, i.e. $\|\hat{X}' - \hat{X}\|_2$, to quantify how anomalous is a sample [1]. We use this principle to establish a surrogate measure of the segmentation quality by quantifying a segmented mask and its reconstruction.

Let us so formalize the function $\rho(\hat{X}, \hat{X}')$, a surrogate measure of the segmentation quality of the mask in the absence of GT. In this context, we denote \hat{X}' a pseudo GT (pGT) since it acts as the reference to measure performance. We present two different scenarios for ρ . First, we propose

$$\rho_1 : \mathcal{C}^{H \times W} \rightarrow \mathbb{R}, \quad (1)$$

which represents the most common setup in autoencoder-based anomaly detection. Due to the generic nature of ρ_1 , well-suited metrics for segmentation quality assessment can be used, such as the DSC or the Hausdorff Distance (HD). Secondly, we propose

$$\rho_2 : \mathcal{C}^{H \times W} \rightarrow \mathbb{R}^{H \times W}. \quad (2)$$

This function generates a visual map of the inconsistencies between the two masks. We use as $\rho_2(\cdot)$ a pixel-wise XOR operation between the segmentation mask \hat{X} and the pGT.

These two types of measures can be used jointly for performance assessment and model monitoring. Measures obtained from ρ_1 -type functions can be paired with a threshold to flag poor segmentation results. The raised alert would then be used to take application-specific countermeasures as, for instance, a visual inspection of an inconsistency map generated by ρ_2 .

Layer	Output Size	Parameters		
		Kernel	Stride	Padding
Input	256x256x4			
Block1	128x128x32	4x4	2	1
Block2	64x64x32	4x4	2	1
Block3	32x32x32	4x4	2	1
Block4	32x32x32	3x3	1	1
Block5	16x16x64	4x4	2	1
Block6	16x16x64	3x3	1	1
Block7	8x8x128	4x4	2	1
Block8	8x8x64	3x3	1	1
Block9	8x8x32	3x3	1	1
Conv2d	4x4x100	4x4	2	1




Fig. 2. Architecture of the encoding module. The decoder is built by reversing this structure and replacing convolutions with transposed convolutions.

Network Architecture. We use the CA architecture proposed in [4] as the backbone network (Fig. 2) with the following modifications. We use a latent space dimension to accommodate 100 feature maps of size 4×4 . A softmax activation function is added to the last layer to normalize the output to a probability distribution over predicted output classes, as well as batch-normalization and dropout to each hidden layer. We use the loss function $\mathcal{L} = \mathcal{L}_{\text{MSE}}(X, X') + \mathcal{L}_{\text{GD}}(X, X')$, where \mathcal{L}_{MSE} is the mean squared error loss and \mathcal{L}_{GD} the generalized dice loss [16]. Trained over 500 epochs, for the first 10 epochs \mathcal{L}_{GD} is computed leaving aside the background class to avoid the convergence to a dummy blank solution. The network weights are set using a *He* normal initializer. The Adam optimizer is initialized with learning rate 2×10^{-4} and a weight decay of 1×10^{-5} . After every epoch, the model is evaluated on the validation set. The weights retrieving the lowest \mathcal{L} value are stored for testing.

3 Experiments and Results

Section 3.1 describes the datasets used, the setup and implementation of the experiments. Experimental results are then presented in section 3.2.

3.1 Experimental Setup

Data. We used data from the MICCAI 2017 Automatic Cardiac Diagnosis Challenge (ACDC) [5]. The dataset consists of an annotated set with 100 short-axis cine magnetic resonance (MR) images, at end-diastole (ED) and end-systole (ES), with corresponding labels for the left ventricle (LV), right ventricle (RV), and myocardium (MYO). The set was split into training and validation subsets using an 80:20 ratio. The challenge also provides a testing set with 50 cases, with no ground truth publicly available. To have uniform image sizes, these were placed in the middle of a 256×256 black square. Those exceeding this size were center cropped.

Setup. We trained the monitoring framework using the ground truth masks from the ACDC training set and used it to assess the performance of five methods participating in the ACDC Challenge [3,6,7,17,19] and an additional state-of-the-art cardiac segmentation model [2]. We trained five models [2,3,6,7,19] using the challenge’s full training set (MR images and masks) and then segmented the ACDC test images. For the remaining method [17], we obtained the segmentation masks directly from the participating team.

The segmentations from every method were fed to the monitoring framework. The resulting pGTs were used to compute ρ_1 -type measures (Def. 1), the DSC and the HD, and a ρ_2 -type measure, an inconsistency map (Def. 2). We also computed pseudo DSC/HD using the RCA [18]. The ACDC challenge platform estimates different performance measures (DSC, HD, and other clinical measures) on the testing set upon submission of the segmentation results. We uploaded the masks from every model to obtain real DSC and HD. To differentiate the real measures computed by the platform from our estimates, we denote the latter ones pDSC and pHD. In our experiments, we set $\text{pHD} > 50$ or $\text{pDSC} < 0.5$ to flag a segmentation as suspicious and $\text{pHD} = \text{pDSC} = 0$ to raise an erroneous segmentation alert flag.

Implementation. We implemented our framework in PyTorch. All the cardiac segmentation models used the available implementations, except for [17] where we had the segmentation masks. The RCA was implemented following the guidelines in [18] using a previously validated atlas propagation heart segmentation framework [22]. All experiments ran on Amazon Web Services with a Tesla T4 GPU. To encourage reproducibility, our code and experiments are publicly available from a Github repository¹.

3.2 Results

Figure 3 presents scatter plots of the real DSC and HD from the ACDC platform and the pDSC and pHD obtained with our framework and the RCA [13,18]. We present results for LV, RV, and MYO over all generated segmentation masks, and report the Pearson correlation coefficient r .

The results show a high positive correlation between real scores and our estimations. Our framework consistently outperforms RCA. For both RCA and our proposed approach, the real and pseudo HDs show stronger correlations than the DSC. This can be explained by the higher sensitivity of the HD to segmentation errors. Instead, the DSC shows little variability in the presence of minor segmentation, suggesting that both methods have difficulties in modeling small deviations from the reference ground truth data, i.e. very high-quality segmentations.

Table 1 simulates the ACDC Challenge results using real HD and pHD for every model to determine if our framework is a reliable means to rank the performance of the different cardiac segmentation methods. We do not use the pDSC

¹ https://github.com/robustml-eurecom/quality_control_CMR

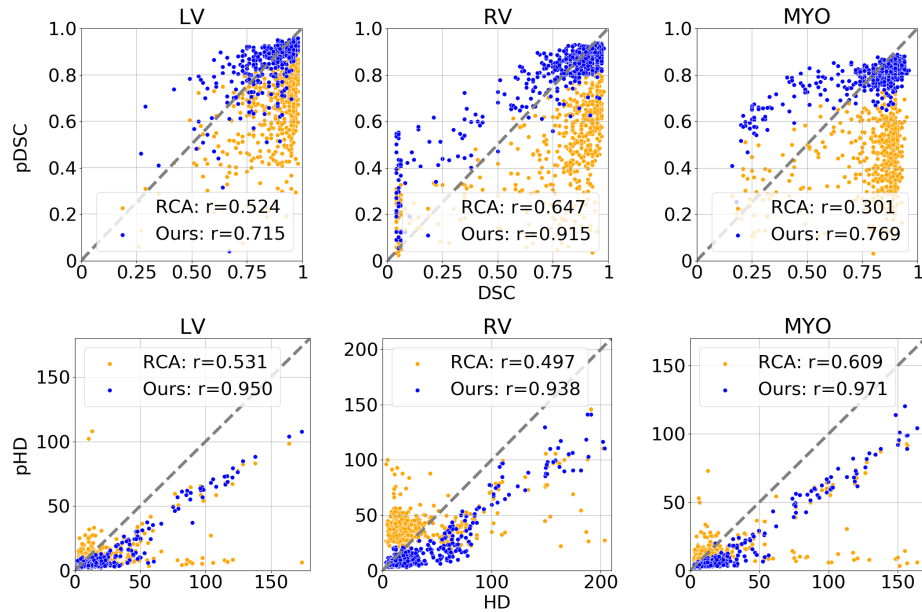


Fig. 3. DSC vs, pDSC (top) and HD vs. pHD (bottom) for our framework (blue dots) and RCA (yellow dots) on the left ventricle (LV), right ventricle (RV) and myocardium (MYO).

for the ranking, as it reported a lower Pearson correlation coefficient r in our first experiment (Fig. 3).

We report results for LV, RV, and MYO in ED and ES and compare them against the RCA. The ranking quality is assessed using Spearman’s rank correlation coefficient r_s between the real and the pseudo measures. The r_s assesses if there is a monotonic relationship between both measures, i.e. it allows to determine if the pseudo measure is a valid criterion to rank the different methods. For fairness in the comparison, the ranking does not include one method [6], where RCA failed. In five out of six cases, we were able to perfectly reproduce the real ranking ($r_s=1.0$). In the remaining case, the left ventricle in end-diastole (ED), there is only one difference between the real and our pseudo ranking ($r_s=0.9$), where the 3rd and 4th places were swapped. The positive high r_s scores obtained by our framework confirm that it is a reliable mean for method ranking.

Through the use of alert flags we were able to detect 16 cases for which the challenge platform had reported NaN values indicating errors in the submitted results. Fifteen cases were flagged as erroneous (pHD = pDSC = 0) and one as suspicious (pHD > 50) by our framework. Fig. 4 middle and bottom row illustrates two of these cases. The middle row presents a segmentation flagged as erroneous, where the inconsistency map confirms that the left ventricle has not been segmented. The bottom row shows the case of a segmentation flagged

Table 1. ACDC Challenge simulation with 6 models using the real HD (GT) and the pHD obtained with our framework (Ours) and the RCA. HD and pHD reported in mm. The Spearman’s rank correlation coefficient r_s measures the ranking accuracy (the closer to 1.0 the better). pHD scores using RCA are excluded for [6], where the registration step failed.

ED									
Model	LV			RV			MYO		
	GT	Ours	RCA	GT	Ours	RCA	GT	Ours	RCA
Bai [2]	39.01	23.38	15.55	50.21	31.82	56.22	47.10	28.46	20.42
Baumgartner [3]	7.14	3.87	9.30	14.00	7.72	37.63	9.49	4.43	10.52
Isensee [6]	7.01	3.88	-	11.40	7.82	-	8.44	4.38	-
Khened [7]	16.81	6.39	10.58	13.25	6.87	39.01	16.09	6.08	11.22
Tziritas [17]	8.90	4.69	8.92	21.02	9.86	41.10	12.59	4.58	10.65
Yang [19]	16.95	5.29	12.96	86.08	47.24	44.75	31.93	16.39	15.12
r_s	-	0.90	0.90	-	1.00	0.80	-	1.00	1.00

ES									
Model	LV			RV			MYO		
	GT	Ours	RCA	GT	Ours	RCA	GT	Ours	RCA
Bai [2]	50.53	29.56	20.01	52.73	31.40	53.68	52.72	31.05	26.60
Baumgartner [3]	10.51	4.41	9.56	16.32	7.10	35.50	12.47	4.77	9.33
Isensee [6]	7.97	4.07	-	12.07	6.99	-	7.95	4.27	-
Khened [7]	20.14	6.96	11.72	14.71	7.07	35.65	16.77	6.03	10.36
Tziritas [17]	11.57	5.00	10.46	25.70	9.61	36.51	14.78	5.59	10.60
Yang [19]	19.13	6.11	11.78	80.42	33.21	40.68	32.54	16.98	13.68
r_s	-	1.00	0.90	-	1.00	0.80	-	1.00	0.90

as suspicious, where the left ventricle’s pHD is high (pHD = 104.93), although the pDSC = 0.814 is within normal range. The inconsistency map confirms the clear segmentation error.

Additionally, Fig. 4 top row illustrates an example of a segmentation mask that has not been flagged, i.e. a segmentation result considered good by the ρ_1 -type metrics of our framework (pHD < 50 and pDSC > 0.5 for all LV, RV and MYO). The inconsistency map (Fig. 4c) flags pixels in the edges of all the structures as suspicious. While it is clear from the images that the segmentation mask for the RV (pDSC = 0.75, DSC = 0.89) has some errors in the edges, which require a revision, the flagged pixels in the LV and MYO are more difficult to assess. In such a case, a more informative ρ_2 -type measure than the selected XOR is desirable.

4 Discussion and Conclusions

We presented a novel learning framework to monitor the performance of cardiac image segmentation models in the absence of ground truth. Our framework addresses the limitations of previous learning-based approaches [8,11,12,14] thanks to its formulation under an anomaly detection paradigm that allows training




Fig. 4. Segmentation masks (a, d and g), pGTs (b, e and h), and inconsistency maps (c, f and i) for **first row:** successful segmentation, according to the alert thresholds; **middle row:** a segmentation mask flagged as erroneous with $pHD = pDSC = 0$ in the LV; and **bottom row:** a segmentation mask flagged as suspicious with $pHD = 104.93$ for the LV ($pDSC = 0.814$). The inconsistency maps confirm the errors.

without requiring quality scores labels. Our results show a good correlation between real performance measures and those estimated with the pseudo ground truth (pGT), making it a reliable alternative when there is no reference to assess a model. Compared with state-of-the-art RCA [13,18], our method avoids the use of image registration which makes it more robust, scalable, and considerably faster (~ 20 min RCA vs. ~ 0.2 s ours, per case). CAs allow for fast inference which conforms to real-time use, thus permitting a quick quality assignment, for example, in a clinical setting. All these characteristics make the proposed

framework a viable option for quality control and system monitoring in clinical setups and large population studies.

A current limitation of the proposed framework is that it is less reliable when assessing high-quality segmentations. A simple and practical way to address this could be to increase the alert thresholds (lower pHDs, higher pDSCs), making them more sensitive to small segmentation errors, and to develop more informative ρ_2 -type measures, similar to those proposed by uncertainty quantification methods [11]. However, we consider that a more principled approach should be favored by investigating mechanisms to model these smaller deviations from the reference.

Another straightforward extension of this work could be to embed the proposed framework within the segmenter network, as a way to perform quality control during training. This idea has been previously explored by different AE-based segmentation methods [9,10,20]. However, the results reported in [10] suggest that this does not fully solve the problem of unexpected erroneous and anatomically impossible results, originally reported in [5]. Therefore, we consider that separate and modular monitoring frameworks should be favored over end-to-end solutions and future works should focus on improving the limitations of the current quality control techniques, in particular, their poor sensitivity to smaller errors.

Acknowledgments. The authors would like to thank Christian Baumgartner, Elios Grinias, Jelmer M. Wolterink, and Clement Zotti for their help in the reproduction of the ACDC Challenge rankings by sharing their code or results and, overall, through their valuable advice.

References

1. Audibert, J., Michiardi, P., Guyard, F., Marti, S., Zuluaga, M.A.: USAD: unsupervised anomaly detection on multivariate time series. In: Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. pp. 3395–3404 (2020)
2. Bai, W., Sinclair, M., Tarroni, G., Oktay, O., Rajchl, M., Vaillant, G., Lee, A.M., Aung, N., Lukaschuk, E., Sanghvi, M.M., Zemrak, F., Fung, K., Paiva, J.M., Carapella, V., Kim, Y.J., Suzuki, H., Kainz, B., Matthews, P.M., Petersen, S.E., Piechnik, S.K., Neubauer, S., Glocker, B., Rueckert, D.: Automated cardiovascular magnetic resonance image analysis with fully convolutional networks. *Journal of Cardiovascular Magnetic Resonance* **20**(1), 65 (2018)
3. Baumgartner, C., Koch, L., Pollefeys, M., Konukoglu, E.: An exploration of 2D and 3D deep learning techniques for cardiac MR image segmentation. In: Statistical Atlases and Computational Models of the Heart. pp. 111–119 (2017)
4. Bergmann, P., Löwe, S., Fauser, M., Sattlegger, D., Steger, C.: Improving unsupervised defect segmentation by applying structural similarity to autoencoders. In: 14th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications. pp. 372–380 (2019)

5. Bernard, O., Lalande, A., Zotti, C., Cervenansky, F., Yang, X., Heng, P.A., Cetin, I., Lekadir, K., Camara, O., Ballester, M.A.G., Sanroma, G., Napel, S., Petersen, S., Tziritas, G., Grinias, E., Khened, M., Kollerathu, V.A., Krishnamurthi, G., Rohe, M.M., Pennec, X., Sermesant, M., Isensee, F., Jager, P., Maier-Hein, K.H., Full, P.M., Wolf, I., Engelhardt, S., Baumgartner, C.F., Koch, L.M., Wolterink, J.M., Isgum, I., Jang, Y., Hong, Y., Patravali, J., Jain, S., Humbert, O., Jodoin, P.M.: Deep learning techniques for automatic MRI cardiac multi-structures segmentation and diagnosis: Is the problem solved? *IEEE Transactions on Medical Imaging* **37**(11), 2514–2525 (2018)
6. Isensee, F., Jaeger, P.F., Full, P.M., Wolf, I., Engelhardt, S., Maier-Hein, K.H.: Automatic cardiac disease assessment on cine-MRI via time-series segmentation and domain specific features. In: *Statistical Atlases and Computational Models of the Heart*. pp. 120–129 (2017)
7. Khened, M., Alex, V., Krishnamurthi, G.: Densely connected fully convolutional network for short-axis cardiac cine MR image segmentation and heart diagnosis using random forest. In: *Statistical Atlases and Computational Models of the Heart*. pp. 140–151 (2017)
8. Kohlberger, T., Singh, V., Alvino, C., Bahlmann, C., Grady, L.: Evaluating segmentation error without ground truth. In: *Medical Image Computing and Computer-Assisted Intervention*. pp. 528–536 (2012)
9. Oktay, O., Ferrante, E., Kamnitsas, K., Heinrich, M., Bai, W., Caballero, J., Cook, S.A., De Marvao, A., Dawes, T., O’Regan, D.P., et al.: Anatomically constrained neural networks (acnns): application to cardiac image enhancement and segmentation. *IEEE Transactions on Medical Imaging* **37**(2), 384–395 (2017)
10. Painchaud, N., Skandarani, Y., Judge, T., Bernard, O., Lalande, A., Jodoin, P.M.: Cardiac segmentation with strong anatomical guarantees. *IEEE Transactions on Medical Imaging* **39**(11), 3703–3713 (2020)
11. Puyol-Antón, E., Ruijsink, B., Baumgartner, C.F., Masci, P.G., Sinclair, M., Konukoglu, E., Razavi, R., King, A.P.: Automated quantification of myocardial tissue characteristics from native t1 mapping using neural networks with uncertainty-based quality-control. *Journal of Cardiovascular Magnetic Resonance* **22**(1) (2020)
12. Robinson, R., Oktay, O., Bai, W., Valindria, V., Sanghvi, M., Aung, N., Paiva, J., Zemrak, F., Fung, K., Lukaschuk, E., Lee, A.M., Carapella, V., Jin Kim, Y., Kainz, B., Piechni, S.K., Neubauer, S., Petersen, S.E., Rueckert, D., Glocker, B.: Real-time prediction of segmentation quality. In: *Medical Image Computing and Computer Assisted Intervention*. vol. 11073, pp. 578–585 (2018)
13. Robinson, R., Valindria, V.V., Bai, W., Oktay, O., Kainz, B., Suzuki, H., Sanghvi, M.M., Aung, N., Paiva, J.M., Zemrak, F., Fung, K., Lukaschuk, E., Lee, A.M., Carapella, V., Kim, Y.J., Piechnik, S.K., Neubauer, S., Petersen, S.E., Page, C., Matthews, P.M., Rueckert, D., Glocker, B.: Automated quality control in image segmentation: application to the UK biobank cardiovascular magnetic resonance imaging study. *Journal of Cardiovascular Magnetic Resonance* **21**(1) (2019)
14. Ruijsink, B., Puyol-Antón, E., Oksuz, I., Sinclair, M., Bai, W., Schnabel, J.A., Razavi, R., King, A.P.: Fully automated, quality-controlled cardiac analysis from CMR. *JACC: Cardiovascular Imaging* **13**(3), 684–695 (2020)
15. Sculley, D., Holt, G., Golovin, D., Davydov, E., Phillips, T., Ebner, D., Chaudhary, V., Young, M., Crespo, J.F., Dennison, D.: Hidden technical debt in machine learning systems. In: *Advances in Neural Information Processing Systems* 28. pp. 2503–2511 (2015)

16. Sudre, C.H., Li, W., Vercauteren, T., Ourselin, S., Cardoso, M.J.: Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. In: Deep learning in medical image analysis and multimodal learning for clinical decision support. pp. 240–248 (2017)
17. Tziritas, G., Grinias, E.: Fast fully-automatic localization of left ventricle and myocardium in MRI using MRF model optimization, substructures tracking and B-spline smoothing. In: Statistical Atlases and Computational Models of the Heart. pp. 91–100 (2017)
18. Valindria, V.V., Lavdas, I., Bai, W., Kamnitsas, K., Aboagye, E.O., Rockall, A.G., Rueckert, D., Glocker, B.: Reverse classification accuracy: Predicting segmentation performance in the absence of ground truth. *IEEE Transactions on Medical Imaging* **36**(8), 1597–1606 (2017)
19. Yang, X., Bian, C., Yu, L., Ni, D., Heng, P.A.: Class-balanced deep neural network for automatic ventricular structure segmentation. In: Statistical Atlases and Computational Models of the Heart. pp. 152–160 (2017)
20. Yue, Q., Luo, X., Ye, Q., Xu, L., Zhuang, X.: Cardiac segmentation from lge mri using deep neural network incorporating shape and spatial priors. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 559–567 (2019)
21. Zuluaga, M.A., Burgos, N., Mendelson, A.F., Taylor, A.M., Ourselin, S.: Voxel-wise atlas rating for computer assisted diagnosis: Application to congenital heart diseases of the great arteries. *Medical Image Analysis* **26**(1), 185–194 (2015)
22. Zuluaga, M.A., Cardoso, M.J., Modat, M., Ourselin, S.: Multi-Atlas Propagation Whole Heart Segmentation from MRI and CTA Using a Local Normalised Correlation Coefficient Criterion. In: Functional Imaging and Modeling of the Heart (FIMH). vol. 7945, pp. 174–181 (2013)